

A Holding Function for Conflict Probe Applications

Dave McNally*

NASA Ames Research Center, Moffett Field, CA 94035

Joe Walton†

University of California, Santa Cruz

NASA Ames Research Center, Moffett Field, CA 94035

Under current National Airspace System operations conflict alerts for aircraft in holding patterns are often missed or in error due to the fact that trajectories for holding aircraft are not modeled in existing Conflict Alert or Conflict Probe automation. In addition, a controller in one sector may not know when aircraft are holding in a neighboring sector. These factors can lead to an increased potential for loss of separation while aircraft are flying in holding patterns. The objective of this work is to develop a holding function that automatically determines when an aircraft enters a holding pattern, computes a holding region around the pattern, and probes the holding region for conflict with other traffic. Since controller workload is generally high during periods when aircraft are in holding the operational concept of use assumes the holding region is automatically computed and the controller is alerted only if another aircraft is predicted to fly through the holding region. The holding algorithm is applied only to aircraft that would likely go into holding during rush periods, e.g., arrivals to capacity constrained airports. A holding region is computed for aircraft that settle out on a steady outbound course. Initial estimates of the holding fix position, turn radius, and holding pattern leg length are computed and automatically updated as the aircraft flies the pattern. The holding algorithm was implemented in the Center/TRACON Automation System software suite and tested using Host radar track and flight plan data from the Fort Worth Air Route Traffic Control Center. Of 37 aircraft that went into holding during a 1 hour severe weather period, the holding function correctly computed and updated 34 holding regions that accurately reflected the holding pattern airspace. One of these aircraft was involved in an operational error that would certainly have been prevented had the 3 min holding alert been displayed to the controller. In 3 cases the holding region was activated incorrectly, then deactivated following subsequent track updates, and later computed correctly on the second outbound leg. The results show that Host track and flight plan data alone may be used to automatically model and conflict probe holding airspace in real-time.

I. Introduction

When the number of aircraft approaching a busy airport exceeds the landing rate capacity of the airport air traffic controllers and traffic managers take steps to meter the flow of arrival traffic into the airport. Arrival metering could be triggered by severe weather, runway closures, unusually heavy arrival demand, or a combination of all these factors. Controllers meter arrival flows through speed control, delay vectors, 360 deg turns, and in the worst case, when delay requirements are particularly large, by placing aircraft in holding patterns. Metering delays of over 5 min typically require aircraft to be placed in holding patterns. When aircraft are flying in holding patterns National Airspace System (NAS) procedures require that controllers protect the holding pattern airspace, that is, ensure that other aircraft do not fly through the holding pattern at the holding altitude.

Conflict detection functions such as Conflict Alert and Conflict Probe do not account for aircraft in holding patterns. Conflict Alert projects forward based on analysis of the last 2-4 radar track measurements without regard for flight plan intent and, as a result, produces erroneous results for aircraft in turning flight. Conflict Probe trajectories merge radar track measurements and flight plan intent information and are therefore able to account for turns at downstream flight plan waypoints. However, NAS flight plans do not contain the information required to

* Principal Investigator, Automation Concepts Research Branch, MS 210-10, Member

† Principal Software Engineer, UC Santa Cruz, MS 210-8

model a holding pattern. Neither Conflict Alert nor Conflict Probe have the ability to model a holding trajectory or to probe a non-holding trajectory for conflict with the holding pattern airspace.

The objective of this work is to develop a real-time function that automatically detects when an aircraft has entered a holding pattern, models the holding pattern airspace region, and then probes the trajectories of other non-holding aircraft for conflict with the holding region. The controller is alerted if an aircraft is predicted to fly through the holding region at the holding altitude. An important aspect of the design is to automatically determine when aircraft have entered a holding pattern and to automatically compute and dynamically update the holding region. The concept and functionality may be extended to allow the controller to manually activate and deactivate the holding function.

The paper is organized as follows. Holding operations in en route Center airspace are described briefly. The next three sections describe the holding methodology, the detailed logic for automatic detection and modeling of the holding pattern airspace, and the holding conflict logic. The performance of the holding pattern modeling and the holding conflict detection using recorded data from the FAA's Fort Worth Air Route Traffic Control Center (ARTCC, or Center) is described in the next two sections. Ideas for a controller's user interface and extensions of the logic to make use of multiple aircraft holding at the same fix are described in the last two sections. The paper closes with some concluding remarks.

II. Holding Operations

A holding pattern is defined by a holding fix, a reference course, a turn direction (left or right) and a leg length. There may be more than one aircraft in the same holding pattern but they all must be separated vertically (by legal standards) regardless of their horizontal position in the pattern. The reference course defines the orientation of the pattern. Leg lengths are nominally 1.5 minutes of flight or 10-20 miles. A standard holding pattern is shown in Fig. 1. The holding fix may be on the aircraft route of flight, or adapted arrival route, or it may be off the route of flight. A controller's holding clearance to the pilot specifies the holding fix, the holding altitude, the outbound (reference) course, the turn direction, and the leg length. Aircraft are instructed to fly first to the holding fix and then turn to enter the holding pattern. Many standard holding patterns are defined on airman's approach plates and specify the holding fix, course, and turn direction. A controller's holding clearance may, therefore, be to fly the standard pattern at a specified fix and altitude.

As an aircraft enters a holding pattern it may approach the holding fix from any direction. Figure 2 shows the three standard approach trajectories by which an aircraft may enter a holding pattern [1]. From our observation of a limited amount of Fort Worth Center data, it appears that it is most common for aircraft to enter the holding pattern from the inbound direction and turn outbound at the fix as in Fig. 2, approach (a). However, the 3 standard approaches should be accounted for in the detection and modeling of the holding pattern.

As arrival aircraft in Center airspace approach the terminal area they fly a course that is either along the standard terminal arrival route (STAR) or possibly direct to the arrival metering fix, where they pass from Center to terminal (TRACON) airspace. When it becomes necessary to meter arrival flows into the airport controllers use other, less extreme, methods for metering before putting aircraft into holding patterns. Metering vectors of up to 90 deg or 360 turns are commonly used to absorb required delay. It is therefore important to distinguish between aircraft on metering vectors, or 360 degree turns, and those entering a holding pattern.




Figure 1. Standard holding pattern.




Figure 2. Standard approaches to a holding pattern

III. Holding Detection and Modeling

A. Methodology

The design objective is to develop an automated methodology for detection and modeling of holding pattern airspace. Since controller workload is generally high during periods when aircraft are holding, holding detection, modeling, and conflict probing should be fully automatic and not require any controller input. Therefore, it is assumed that the only input data available are Host radar track messages, which update every 12 sec, and Host flight plans.

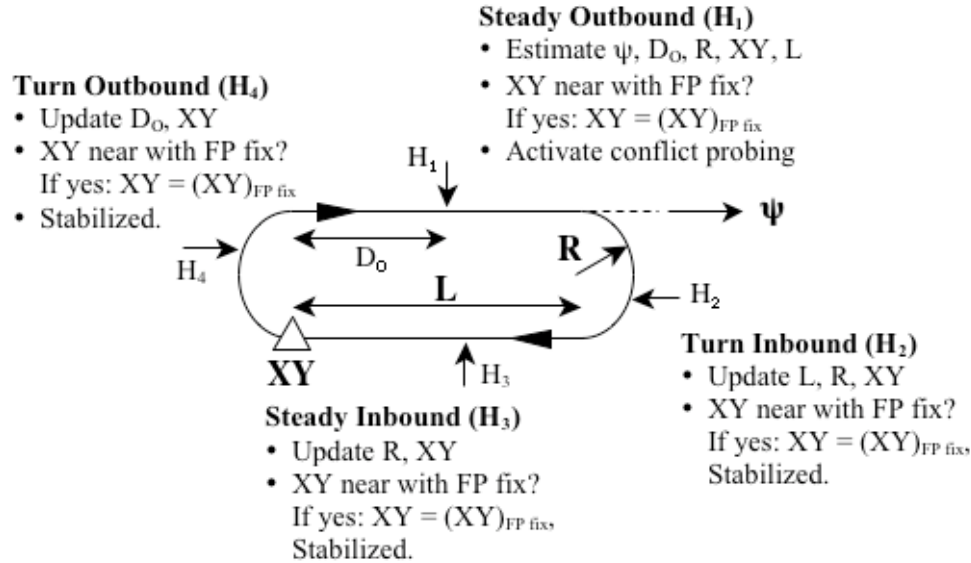


Figure 3. Automated holding detection and modeling methodology

The automatic holding detection and modeling methodology is summarized in Fig. 3. Any holding pattern is fully specified by the xy coordinates of the holding fix, the turn direction (left or right), the outbound course, ψ , the turn radius, R , and the leg length, L . These quantities must be estimated in real-time in order to model the holding pattern. The conflict detection approach is to probe the trajectories of non-holding aircraft for conflict with a rectangular volume that surrounds the holding pattern and is vertically centered at the altitude of the holding aircraft. These holding volumes are treated like special use airspace in terms of how the trajectories of non-holding aircraft are probed for conflict with the holding airspace. Even though it is common for multiple aircraft to hold at the same fix, but at different altitudes, in this formulation, the holding regions for different aircraft in the same pattern are computed independent of one another.

A primary assumption in this methodology is that an aircraft is entering a holding pattern if it is destined for, and close to, an airport that could be capacity constrained (e.g., major airports like the Dallas/Fort Worth International Airport (DFW)), and it turns and settles out on a steady outbound course. In order to prevent aircraft on arrival metering vectors from being identified as holding aircraft, the steady outbound course must be at least 110 deg away from a direct-to-next-fix course or a direct-to-meter-fix course.

As shown in Fig. 3, the a holding region is initialized at the point H_1 where the aircraft settles out on a steady outbound course. As the aircraft flies the pattern improved estimates of L , R , and XY are used to improve and update the rectangular holding region. The estimated course, ψ , at the H_1 point defines the orientation of the pattern. The turn radius, R , is estimated using ground speed and an assumed bank angle of 25 deg. We have observed that it usually takes about 6-8 track updates for an aircraft to settle out on a steady course following a turn outbound. Therefore, for our initial estimate of the distance D_0 we assume 7 track updates from the point where the aircraft is turning outbound, and is adjacent to the holding fix, to the point where the aircraft settles out on a steady outbound course. The xy coordinates of the holding fix are estimated using the track position at H_1 and our estimates of ψ , R and D_0 . If the estimated position of the holding fix is within 3 nmi of any flight plan fix, then the xy coordinates of the holding fix are replaced with the xy coordinates of the correlated flight plan fix. The holding region is initialized assuming $L=10$ nmi and a 5 nmi buffer on all sides of the pattern.

The leg length, turn radius and holding fix estimates are updated at the point H_2 where the aircraft turns inbound. The maximum outbound distance at H_2 fixes the leg length. The radius update is computed assuming a steady turn

rate at the H2 point. If the estimated holding fix position correlates with (i.e., is within 3 nmi of) a flight plan fix at the H2 point, the holding pattern is completely specified and the holding model is complete. If a flight plan fix is not correlated at the H2 point, the turn radius is again updated at the point H3 where the aircraft settles out on a steady inbound course. The holding fix estimate is again updated and checked for correlation with a flight plan fix. Finally, if the holding fix was not correlated at the H3 point, the final update to the pattern occurs at the H4 point where the aircraft turns outbound for the second time. Assuming a good estimate of the turn radius at H3, the remaining uncertainty in the holding fix position is along the direction of the reference course, ψ , and may all be attributed to an error in D_O . The holding model is complete at the H4 point for any aircraft whos holding fix was not correlated at the point points H1, H2, or H3. It should be noted that the holding model may stabilize at the H2 point for aircraft holding at a flight plan fix, but will always stabilize at the H4 point for aircraft holding at an off-route fix. The remainder of Section III describes the holding logic in more detail including the logic that determines when to deactivate the holding model.

B. Holding Detection

In order to automatically compute a holding region it is necessary to define observable information that indicates when an aircraft is likely entering a holding pattern. In most cases aircraft entering a holding pattern are near their destination airport. In this analysis only aircraft that are arrivals to either DFW or Dallas Love Field (DAL) are tested to determine if they have entered a holding pattern. DFW and DAL are the major airports in ZFW airspace and aircraft having to enter holding patterns would most likely be arrivals to these airports.

Figure 4 is a sketch of a typical time history of radar track (dots) and estimated course (arrows) for an aircraft entering a holding pattern. The figure also shows the flight plan arrival route (dashed line), and 5 flight plan fixes (triangles). The aircraft is holding at one of the fixes and the last fix on the route is the meter fix. The distinguishing characteristic of the trajectory in Fig. 4 is that the aircraft has made a large turn and then settled out on a steady outbound course. The steady course is not capturing any flight plan waypoint and is not direct to the meter fix. The observation of a large turn followed by a steady outbound course is the basis for initialization of the holding region.

Aircraft course and turn status are computed through analysis of the Center Host radar track messages. A tracking filter computes aircraft course, ψ , and turn status (left, right, or not-turning) at each 12 sec track update [2]. The tracking filter is based on a best fit of the last 6 radar track measurements.

The first step in the holding analysis is to define variables that will allow the holding algorithm to distinguish between aircraft on their flight plan or those on metering vectors, with those that might be entering a holding pattern. $\Delta\psi_{NF}$ is defined as the difference between the current course (ψ) and a course that is direct to the next downstream flight plan (or STAR) fix (ψ_{NF}). $\Delta\psi_{MF}$ is defined as the difference between the current course and a course that is direct to the arrival meter fix (ψ_{MF}). Both of these variables, defined in Eq. (1), are computed and monitored at every track update for all aircraft that qualify for the holding analysis (e.g., all DFW and DAL arrivals).




Figure 4. Typical track and estimated course history for aircraft entering a holding pattern.

$$\begin{aligned}\Delta\psi_{NF} &= |\psi - \psi_{NF}| \\ \Delta\psi_{MF} &= |\psi - \psi_{MF}|\end{aligned}\tag{1}$$

The turn status, which is output from the tracking filter [2] at every track update, is analyzed to determine the turn direction of holding pattern, i.e., left or right, and the point at which the aircraft has settled out on a steady course. It is known from holding operations that an aircraft flies through a relatively large turn as it enters a holding pattern. The turn status data are continuously monitored to detect the direction of the most recent steady turn with a magnitude of 60 deg or greater. A steady turn is one where the turn status remains constant at either right or left for the required amount of course change (60 deg). The requirement for a large steady turn ensures that the turn direction is correctly computed prior to initializing the holding region. For example, when the turn status changes from not-turning to right-turn (as in Fig. 4), the turn status is monitored until it changes back to not turning. The

magnitude of the turn is the difference between the course at the beginning and end points of the turn. The holding region is initialized when the following 3 requirements are simultaneously satisfied at a single track point:

- 1) $\Delta\psi_{NF}$ and $\Delta\psi_{MF}$ are both greater than 110 deg,
- 2) aircraft is on a steady course (turn status = not turning) for 3 consecutive track updates,
- 3) aircraft executed a turn of greater than 60 deg prior to steady course state.

One-hundred-ten (110) deg was selected as the criteria for $\Delta\psi_{NF}$ and $\Delta\psi_{MF}$ primarily to ensure that holding regions are not activated for aircraft on metering vectors. Metering vectors, which are commonly used when controllers are metering arrival flow, are typically no greater than 90 deg.

At the track point where these holding requirements are satisfied the holding status H is set equal to 1 ($H=1$), the holding region is initialized (described in next section) and probing of the holding region is activated (described in a later section). We define this as the initialization point ($H=1$ point). The aircraft position coordinates (x_1, y_1), ground speed (V_{g1}), and course (ψ_1) are stored at the $H=1$ point for later analysis.

C. Holding Airspace Model

The holding region is modeled as a rectangle that surrounds the holding pattern as shown in Fig 5. Estimates of the outbound course, turn direction, turn radius, holding fix, and the holding leg length are required to compute the holding region. The outbound course is set equal to the course at the initialization point ($\psi_H = \psi_1$). The turn direction is the direction of the large turn that preceded the $H=1$ point.

The offset D_O is defined as the distance the aircraft has flown in the outbound direction since it was adjacent to the holding fix (see Fig. 5). From our analysis of the tracking filter transient response for many different aircraft in holding we have observed that it usually takes about 6-8 track updates (beyond the point adjacent to the holding fix) for estimated course to settle out from a turning status to a not-turning status (see Fig. 4). For the purpose of estimating the initial offset we assume 7 track updates from the abeam point to the $H=1$ point. The initial estimate of the offset is defined in Eq. (2) where V_{g1} is the measured ground speed at $H=1$ and ΔT is the 12 sec radar track update rate. The initial estimate of the turn radius is defined in Eq. (3). Since the bank angle ϕ is not known at the $H=1$ point, we assume a nominal bank angle ($\phi = 25$ deg) in order to compute our initial estimate of the turn radius.




Figure 5. Holding region model (for a right turn)

$$D_{O1} = V_{g1}(7\Delta T) \quad (2)$$

$$R_1 = V_{g1}^2 / g \tan \phi \quad (3)$$

Setting $D_O = D_{O1}$ and $R = R_1$ an estimate of the xy position of the holding fix (x_{FX}, y_{FX}) is defined in Eqs (4).

for a right turn:

$$\begin{aligned} x_{FX} &= x_1 - D_O \sin \psi_H + 2R \cos \psi_H \\ y_{FX} &= y_1 - D_O \cos \psi_H - 2R \sin \psi_H \end{aligned} \quad (4a)$$

for a left turn:

$$\begin{aligned} x_{FX} &= x_1 - D_O \sin \psi_H - 2R \cos \psi_H \\ y_{FX} &= y_1 - D_O \cos \psi_H + 2R \sin \psi_H \end{aligned} \quad (4b)$$

It is common, though not required, for aircraft to hold at a fix that is on their flight plan arrival route. Since the xy position of all fixes on the flight plan arrival route is known from Center adaptation data, the estimated position of the holding fix (x_{FX}, y_{FX}) is compared to the known position of all arrival route fixes. If the estimated holding fix

is within a parameter distance from any arrival route fix (3-letter fix or 5-letter intersection), then the holding fix coordinates are replaced with the coordinates of the arrival route fix. A parameter distance of 3 nmi was used in this study. If the position of the holding fix is known, the offset distance D_O and the turn radius R are computed directly using the position coordinates x_1 and y_1 , the reference course ψ_1 , and the xy position of the holding fix.

Referring to Fig. 5, the corner points that define the holding region are defined in Eqs. (5). Since holding patterns typically have leg lengths of 10 to 20 nmi, we set the initial leg length to 10 nmi ($L = L_1 = 10$ nmi). As described in Section IIID, the leg length will be updated as the aircraft flies outbound. The buffer is set to a constant value of $B = 5$ nmi. Note that the formulations in Eqs. (5a) and (5b) are based on the point h1 being closest to the holding fix for both a left and a right turn. The point h2 is adjacent to h1 on the fix-end boundary; h4 is adjacent to h1 on the non-holding side boundary; h3 is adjacent to h2 on the holding side boundary.

for a right turn:

$$\begin{aligned} x_{h1} &= x_{FX} - R\sin\psi_H - B\sin\psi_H + B\cos\psi_H \\ y_{h1} &= y_{FX} - R\cos\psi_H - B\cos\psi_H - B\sin\psi_H \\ x_{h2} &= x_{h1} - (2R + 2B)\cos\psi_H \\ y_{h2} &= y_{h1} + (2R + 2B)\sin\psi_H \\ x_{h3} &= x_{h2} + (2R + 2B + L)\sin\psi_H \\ y_{h3} &= y_{h2} + (2R + 2B + L)\cos\psi_H \\ x_{h4} &= x_{h1} + (2R + 2B + L)\sin\psi_H \\ y_{h4} &= y_{h1} + (2R + 2B + L)\cos\psi_H \end{aligned} \quad (5a)$$

for a left turn:

$$\begin{aligned} x_{h1} &= x_{FX} - R\sin\psi_H - B\sin\psi_H - B\cos\psi_H \\ y_{h1} &= y_{FX} - R\cos\psi_H - B\cos\psi_H + B\sin\psi_H \\ x_{h2} &= x_{h1} + (2R + 2B)\cos\psi_H \\ y_{h2} &= y_{h1} - (2R + 2B)\sin\psi_H \\ x_{h3} &= x_{h2} + (2R + 2B + L)\sin\psi_H \\ y_{h3} &= y_{h2} + (2R + 2B + L)\cos\psi_H \\ x_{h4} &= x_{h1} + (2R + 2B + L)\sin\psi_H \\ y_{h4} &= y_{h1} + (2R + 2B + L)\cos\psi_H \end{aligned} \quad (5b)$$

D. Outbound Holding Region Updates

Once the holding region has been initialized (at $H=1$) the track data are analyzed as the aircraft flies outbound to update the holding leg length. Figure 6 shows the model used to update the leg length and the progression of the radar track data beyond the $H=1$ point. The estimated maximum distance the aircraft will fly in the outbound direction beyond the $H=1$ point is expressed in Eq. (6). The actual distance the aircraft has flown in the outbound direction at each track point beyond the $H=1$ point is expressed in Eq. (7). We define the holding state $H=2$ as the point where the aircraft has flown the maximum distance outbound and started to turn back inbound. In order to determine the $H=2$ point two new variables are defined. ΔD_{OB} , defined in Eq. (8), is the difference between Eq. (6) and Eq. (7) at each point beyond $H=1$. δD_{OB} , defined in Eq. (9), is the incremental change in D_{OB} between successive track updates beyond the $H=1$ point.

$$D_{OB1} = L_1 - D_O + R_1 \quad (6)$$

$$D_{OB} = [(x-x_1)^2 + (y-y_1)^2]^{1/2} \cos\theta_{OB} \quad (7)$$

$$\Delta D_{OB} = D_{OB} - D_{OB1} \quad (8)$$

$$\delta D_{OB} = D_{OBi} - D_{OBi-1} \quad (9)$$

At each track update beyond the $H=1$ point the track data are checked to determine if the aircraft has reached a maximum outbound distance and started to fly inbound. We know that δD_{OB} will continuously increase (by about 1 nmi per track update for an aircraft flying 350 kt ground speed) as the aircraft flies outbound. But we must guard against radar track noise which occasionally indicates that the aircraft is moving in the opposite direction of its actual velocity vector. Note from Fig. 6 that at the point where the aircraft is starting to fly inbound




Figure 6. Outbound leg length update model

the lateral offset from the outbound direction of flight is roughly equal to the turn radius R . Prior to the point where the aircraft starts to make the 180 turn, the lateral offset is near zero. In this analysis we have arbitrarily required that the lateral offset at the turn-back point be greater than 40% of the initial estimate of the turn radius R_1 . Therefore, when both of the conditions defined in Eqs. (10) are met at a single track update, the aircraft is said to have reached a maximum outbound distance and is turning inbound. At this point we set $H=2$.

$$\begin{aligned} \delta D_{OB} &< 0.2 \text{ nmi} \\ C_{OB} &> 0.4R_1 \text{ nmi, where } C_{OB} = [(x-x_1)^2 + (y-y_1)^2]^{1/2} \sin\theta_{OB} \end{aligned} \quad (10)$$

The requirement that C_{OB} be greater than 40% of the initial estimate of the turn radius R_1 is a simple way of preventing radar track noise on the outbound leg from causing a false detection of the $H=2$ point.

The leg length is increased dynamically as the aircraft flies outbound from the $H=1$ point and is fixed at the point where the aircraft reaches the turn-back point ($H=2$). In order to update the leg length we examine the sign and magnitude of ΔD_{OB} (Eq. 8) at each track update while the aircraft flies outbound. The leg length update logic is defined as function of the following 3 ranges of ΔD_{OB} :

- 1) If $\Delta D_{OB} < 0.5$ nmi, the aircraft has flown less than the estimated maximum outbound distance. Under this condition the leg length is updated only once at the point where the aircraft has reached a maximum outbound distance and started to fly inbound. This condition corresponds to the $H=2$ point where the updated leg length is defined as $L_2 = L_1 + \Delta D_{OB}$.
- 2) If $0.5 \text{ nmi} < \Delta D_{OB} < 15$ nmi, the leg length is updated at each track update until the aircraft has reached a maximum outbound distance and started to fly inbound. Prior to reaching the maximum outbound distance $L = L_1 + \Delta D_{OB}$. Once the maximum outbound distance is reached ($H=2$), L_2 is defined as in condition 1).
- 3) If $\Delta D_{OB} > 15$ nmi, and the aircraft has not yet reached a maximum outbound distance and started to fly inbound, the holding region (and probing of the region) are deactivated and the holding status is set back to $H=0$. This check limits the length of the holding pattern leg length to 25 nmi and deactivates false holding regions that may be computed for aircraft that are flying a long outbound distance for some reason.

At each point where the leg length is updated the holding region is recomputed using Eqs. (5) with the new value of L . At any point where the holding status is set back to $H=0$, e.g., condition 3) above, the holding detection logic as defined in Section IIIB is re-started completely and the 3 requirements for holding must again be satisfied before a holding region is initialized.

The turn radius R is updated at the $H=2$ point using turn rate information output from the tracking filter [2]. It is assumed that at the $H=2$ point the aircraft is in a reasonably steady turn and, therefore, turn rate information should provide an improved estimate of turn radius (compared to our guess of $\phi = 25$ deg that was used to compute R at the $H=1$ point). Eq. (11) is used to estimate the bank angle at $H=2$ where Ω is the turn rate at $H=2$ and V_{g2} is the ground speed at $H=2$. The bank angle from Eq. (11) and V_{g2} are then used in Eq. (3) to obtain an improved estimate of turn radius (R_2).

$$\phi_2 = \tan^{-1} [V_{g2} \Omega / g] \quad (11)$$

If the requirements for $H=2$ are never satisfied the aircraft will have likely flown out of the holding region which would deactivate the holding region and associated conflict analysis. The logic which handles holding region boundary crossings is described later in Section IIIG.

E. Inbound Turn Radius Update

Once the aircraft has completed the 180 deg turn at the outbound end of the pattern a final update to the turn radius is computed. Following the $H=2$ point, the track and course data are analyzed to determine the point at which the aircraft has settled out on a steady inbound course that is within ± 10 deg of a course that is 180 deg opposite to ψ_1 for 3 consecutive track updates. This point is defined as $H=3$ where it is assumed that the aircraft has settled out on an inbound course and is flying directly towards the holding fix.

Following the turn back point ($H=2$) the aircraft might settle out on a steady course that is not opposite to the outbound course. The aircraft could be departing the holding pattern and rejoining the arrival route or it may be flying a steady course for a short period prior to capturing the inbound leg towards the holding fix. It is possible that noise in the track data has produced a false, or short-duration, indication of a steady course state. The steady course

state is monitored to see if it persists. If a steady course that is well off the expected opposite course persists, or if the aircraft turns in the opposite direction of the previously identified turn direction, or if it flies out of the holding region, then the holding region is deactivated.

Once it has been confirmed that the aircraft is on a steady inbound course (H=3 point) the turn radius is updated. As illustrated in Fig. 7 the perpendicular distance from the H=3 point and the line defined by the point x_1, y_1 and the reference course ψ_1 is equal to twice the actual turn radius. As shown in Fig. 7 the line passing through x_1, y_1 at course ψ_1 is defined by

$$\begin{aligned} y &= mx + y_{\text{int}}, \text{ where} \\ m &= \cot\psi_1 \\ y_{\text{int}} &= y_1 - mx_1. \end{aligned} \quad (12)$$

The perpendicular distance between the track point at H=3 and the line defined by Eq. (12) may be computed directly. The corrected turn radius R_3 is one-half of this distance. The position of the holding fix (Eqs. (4)) and the holding region (Eqs. (5)) are recomputed with R set equal to R_3 . Note that this solution contains a singularity at course angles of 0 deg and 180 deg. In cases where ψ_1 is within 10 deg of 0 or 180 deg, the computed value of $2R_3$ is set equal to the absolute value of the difference between x_1 and x_3 . With a new value of the turn radius ($R=R_3$) the estimated position of the holding fix (Eqs. (4)) is updated and tested for correlation with a flight plan fix as described earlier. The last step in the H=3 update is to update the holding region coordinates (Eqs. (5)) using the updated value of the holding fix, the turn radius, and the leg length that was computed at the H=2 point ($L=L_2$).

F. Offset Correction

If following the H=3 point the holding fix has not been correlated with a flight plan fix, then there remains uncertainty in the position of the holding fix. Since the turn radius was updated at H=3 most of the remaining uncertainty in the holding fix position is along the direction of the reference course. All of this uncertainty is attributed to an error in the offset that was estimated at the H=1 point. As the aircraft flies inbound from the H=3 point the track and course data are analyzed to obtain an improved estimate of the offset, D_O . A model for the offset update is shown in Fig. 8. The formulation is similar to that used for the outbound leg length correction except it first must be determined if at the H=3 point the aircraft is flying towards the estimated holding fix (as shown in Fig. 8) or away from the estimated holding fix. It is most likely that the aircraft is flying towards the holding fix at H=3, but an error in the offset computed at H=1 or a long inbound leg before H=3 could result in the estimated holding fix being behind the aircraft at the H=3 point. This check is done once at the H=3 point using the aircraft course ψ_3 and the xy coordinates of estimated holding fix position.




Figure 7. Inbound turn radius correction model.




Figure 8. Offset correction model

The estimated maximum distance the aircraft will fly inbound from the H=3 point is expressed in Eqs. (13). The actual distance the aircraft has flown inbound at any track point beyond H=3 is expressed in Eq. (14).

$$\text{If aircraft flying towards estimated fix at H=3: } D_{IB3} = [(x_{FX} - x_3)^2 + (y_{FX} - y_3)^2]^{1/2} + R_3 \quad (13a)$$

$$\text{If aircraft flying away from estimated fix at H=3: } D_{IB3} = R_3 - [(x_{FX} - x_3)^2 + (y_{FX} - y_3)^2]^{1/2} \quad (13b)$$

$$D_{IB} = [(x - x_3)^2 + (y - y_3)^2]^{1/2} \cos\theta_{IB} \quad (14)$$

At each track point beyond H=3 the track data are tested to determine if the aircraft has reached a maximum inbound distance and started to fly outbound. This analysis is similar to the outbound analysis defined earlier by Eqs. (10). When the conditions defined by Eqs. (15) are satisfied at a single track update the aircraft is assumed to have reached the maximum inbound distance and turned outbound. We define this point as H=4.

$$\begin{aligned} \delta D_{IB} &< 0.2 \text{ nmi} \\ C_{IB} &> 0.5R_3 \text{ nmi, where } C_{IB} = [(x-x_3)^2 + (y-y_3)^2]^{1/2} \sin\theta_{IB} \end{aligned} \quad (15)$$

Since the H=3 estimate of the turn radius (R_3) is better than the H=1 estimate of the turn radius (R_1), the scale factor used in the inbound lateral offset computation (C_{IB}) may be larger than that used in the outbound lateral offset (C_{OB}). We have arbitrarily set this scale factor to 0.5 as shown in Eq. (15).

At the H=4 point the offset distance D_O may be corrected. As shown in Fig. 8, the difference between the actual inbound distance measured (Eq. (14)) at the H=4 point and the estimated inbound distance that was computed at the H=3 point (Eq. (13)) is equal to the magnitude of the offset correction and is expressed in Eq. (16). Recall that D_O was estimated at the H=1 point using Eq. (2). Since a better estimate of the turn radius was computed at H=3 and the aircraft is close to its maximum inbound position at the H=4 point, a better estimate of D_O may be expressed by Eq. (17). From Fig. 8, the corrected value of D_O will move the holding fix position mostly along the direction of the reference course. In order to maintain the outbound holding region boundary in the same position that resulted from the H=2 leg length correction, the leg length must be corrected in an equal and opposite amount to that of the offset distance. The leg length correction at H=4 is defined in Eq. (18).

$$\Delta D_{IB-MAX} = D_{IB-MAX} - D_{IB3} \quad (16)$$

$$D_{O4} = D_{O1} + \Delta D_{IB-MAX} \quad (17)$$

$$L_4 = L_2 + \Delta D_{IB-MAX} \quad (18)$$

Finally, the position of the holding fix is updated using Eqs. (4) with the updated offset D_{O4} and turn radius R_3 . The holding fix is checked for correlation with a flight plan fix and the holding region is updated using Eqs. (5) and the updated leg length L_4 and turn radius R_3 .

G. Holding Region Boundary Crossings Tests

If the aircraft flies outside of an active holding region we must determine if the holding region should be deactivated or remain active. The holding region must certainly be deactivated when the aircraft flies out because it has been cleared back onto the arrival route to the airport. However, we may not want to deactivate holding if it appears that aircraft is still in holding, but has flown outside the region due to a holding parameter (e.g., turn radius) error that could be corrected at the next update. In this section we define the logic that determines the action taken when the holding aircraft crosses one of the four holding region boundaries. The logic is defined independently for each of the four boundaries. Referring to Fig. 5, the four holding region boundaries are the holding side (h_2h_3), the outbound end (h_3h_4), the non-holding side (h_4h_1), and the inbound end (h_1h_2).

If the aircraft crosses the holding side boundary at any time, the holding region is deactivated ($H=0$). The aircraft could cross the holding side boundary for a number of reasons, all of which suggest that the holding region should be deactivated. The aircraft could be exiting the holding pattern. The initial turn direction estimate could have been incorrect or the reference course ψ_1 could have been substantially in error causing the holding region to be mis-oriented. The holding region could have been mistakenly generated for an aircraft that is not in holding. As will be shown later in the data analysis section crossing a holding side boundary is relatively uncommon and the ψ_1 estimates are usually reasonably accurate.

As described in Section IIID the outbound leg length will increase at each track update if the actual leg length is longer than the initial estimate of $L=10$ nmi. This, in effect, prevents the aircraft from crossing the outbound boundary on the first outbound leg. Recall, however, that if the outbound leg length updates indicate a leg length of more than 25 miles the holding region will be deactivated. Under the current holding logic, if the aircraft crosses the outbound boundary on the second outbound leg the holding region will be deactivated. The logic could be extended to make an additional correction to the leg length on the second outbound leg. Sometimes pilots ask to extend their leg length after the first pass through the pattern.

If the aircraft crosses the non-holding side boundary before H=2 is set, the holding region is deactivated. This is often the case when an aircraft is departing an off-route holding fix to rejoin the arrival route. If the aircraft crosses the non-holding side boundary after H=2 has been set, the course is tested to determine if the aircraft appears to be

capturing the inbound holding pattern course or departing the holding pattern and rejoining the arrival route. If the course is within 45 deg of the expected inbound course ($\psi_1 - 180$ deg) then we increase the turn radius by 25% and re-compute the holding fix and the holding region. At this point we expect that if the aircraft is in holding the aircraft will soon capture an inbound steady course ($H=3$) and the turn radius will be adjusted as defined in Section III E. If following the 25% radius increase, the aircraft crosses the non-holding side boundary a second time before satisfying the requirements for $H=3$, the holding region is deactivated.

A holding aircraft will cross the inbound boundary for one of two reasons. Either the aircraft is departing the holding pattern and flying in towards the airport, or our initial estimate of the offset, and therefore the holding fix position, was in error. If the estimated holding fix has been correlated with a flight plan fix then we have a lot of confidence in the fix position and the turn radius. If this is the case, then, if the aircraft crosses the inbound boundary, the holding region is deactivated. If the holding fix has not been correlated with a flight plan fix and the aircraft crosses the inbound boundary the course is tested to determine if the aircraft appears to be flying in or turning to fly a second outbound leg.

IV. Holding Region Conflict Probe

The conflict logic is relatively straight forward. Equations 5 define the horizontal boundaries of the holding region. The vertical boundaries are defined by the Mode C (pressure) altitude of the holding aircraft and the standard vertical separation requirements for Center airspace. Together the horizontal and vertical boundaries define the 3-dimensional holding volume against which other aircraft trajectories are probed for conflict.

The vertical separation criteria for Center airspace is 1000 ft for aircraft at or below flight level 290 (FL290) and 2000 ft for aircraft above FL290. A 200 ft vertical buffer is included to account for uncertainty in the Mode C altitude measurements received from the aircraft. This 200 ft buffer is commonly used in Conflict Alert and Conflict Probe applications in Center airspace. The minimum and maximum vertical bounds of the holding region are then defined as:

```

Mode_C_Alt <= FL280:
    minimum_alt = Mode_C_Alt - 800 ft
    maximum_alt = Mode_C_Alt + 800 ft
Mode_C_Alt = FL290:
    minimum_alt = Mode_C_Alt - 800 ft
    maximum_alt = Mode_C_Alt + 1800 ft
Mode_C_Alt > FL290:
    minimum_alt = Mode_C_Alt - 1800 ft
    maximum_alt = Mode_C_Alt + 1800 ft

```

V. Analysis with Fort Worth Center Data

A. Holding Detection and Modeling

The holding function has been implemented as part of the en route conflict prediction software in the Center/TRACON Automation System (CTAS) software suite [3,4]. Radar track and flight plan data recorded from the Host computer at the FAA's Fort Worth Air Route Traffic Control Center (ARTCC, or Center) have been used to test and validate the performance of the holding function. Severe weather in the Fort Worth Center airspace caused a large number of aircraft to be put into holding patterns on September 11, 2003. One-hundred and seventeen (117) arrivals to DFW or DAL were in ZFW airspace during the period from 10:32 AM to 11:33 AM local time. The holding function was applied to all 117 arrivals during this period. In this section we show a number of sample holding cases to illustrate the operation of the holding function and its performance under a variety of holding scenarios. The results of the analysis are summarized in Table 1. As shown in the table, the holding pattern was correctly computed for 34 out of 37 cases where aircraft were actually flying holding patterns. In the other 3 cases the holding region was initially computed incorrectly and then corrected. There was 1 case where a holding region was incorrectly computed for a period and then deactivated for aircraft that were flying circular trajectories, albeit ones with periods of steady outbound flight, that triggered activation of a holding region.

Table 1. Summary of holding function performance using Fort Worth Center Host data from Sept 11, 2003 between 10:32 AM and 11:33 AM local time.

Arrival aircraft analyzed for holding	117
Aircraft actually in holding	37
Correctly computed holding regions, holding fix identified	19
Correctly computed holding regions, off-route holding fix or fix not identified	15
Initialized incorrectly due to holding model errors (long outbound prior to fix crossing, incorrect turn radius, or reference course error)	3
False activation for aircraft not in holding (flying circular trajectory with steady outbound course)	1

Figure 9 is an xy-plot of the radar track position and the holding region for an aircraft (AAL1728) holding at the fix CADES. The estimated course and turn status (L, R, or no mark for not turning) are shown at each of the 12-sec radar track updates. The circled track point marked with an “S” is the point where the aircraft first satisfied the delta-course requirements (Eqs. (1)). The estimated holding fix, computed at the H=1 point, is shown (triangle with 1 inside) along with the actual holding fix (dark triangle). At the initialization point (H=1), the flight plan fix CADES was within 3 nmi of the estimated holding fix and was, therefore, automatically identified as the actual holding fix. Since the actual holding fix was known, the offset and the turn radius were computed directly, and the only update to the holding region was the leg length update at the turn-back point (H=2). The dashed rectangle is the initial holding region with a leg length of 10 nmi and the solid rectangle is the final holding region update with a leg length of 15.2 nmi. The holding region is deactivated and the holding analysis is reinitialized (H=0) at the track point where the aircraft exits the holding region.




Figure 9. Radar track and holding region for aircraft holding at fix CADES. Flight plan fix identified at initialization point (H=1).

Figure 10 shows a case (AAL2308) where the estimated holding fix was not correlated with a flight plan fix until the second turn radius update which occurred at the steady inbound point (H=3). The turn radius estimate at the initialization point (H=1) underestimates the actual turn radius and, as a result, the width of the initial holding region (small dashed rectangle). The turn radius update at the turn-back point is more accurate, but still does not correlate with a flight plan fix. Note that if the search radius for correlating an estimated fix with a flight plan fix were increased slightly, say to 5 nmi, the actual holding fix would have been correlated at the turn-back point. The magnitude of the search radius could be determined by the minimum distance between arrival route fixes (3 letter fixes or 5 letter intersections) for the adapted airport.




Figure 10. Holding fix correlated with flight plan fix (CADES) at inbound turn radius update.

Figure 11 shows a case where an aircraft (AAL802) is holding at a fix that is not on the flight plan route. The first two estimates of the turn radius, computed at the initialization point (H=1) and the turn-back point (H=2) are

reasonably accurate, but the initial estimate of the offset is off by about 5 nmi. In this case the offset is corrected at the point ($H=4$) where the aircraft turns back in the outbound direction. The final estimate of the holding fix position appears accurate to within about 1 nmi as indicated by its position relative to the track data.




Figure 11. Aircraft holding at off-route fix showing leg length ($H=2$), turn radius ($H=3$) and offset ($H=4$) corrections and deactivation point ($H=0$).




Figure 12. Aircraft flying circles showing that no holding region is created.

Figure 12 shows an aircraft (CAA803) that meets the course deviation requirements defined in Eq. (1) but never settles out on a steady outbound course so a holding region is never computed.

Aircraft entering a holding pattern will occasionally fly a steady outbound course for a short period and then adjust course to capture the assigned course for the holding pattern. Under these circumstances, the holding region could be initialized with a reference course that is in error. An example of this is shown in Fig. 13 (AAL2416). As is evident in the figure, if the aircraft settles out on a steady outbound course and then adjusts course, the holding region could be skewed with respect to the actual holding pattern. In this example the aircraft flies out of the holding side boundary on the second outbound leg as a result of the reference course error. This deactivates the




Figure 13. AAL2416 holding region with reference course error due to course change on outbound leg.




Figure 14. AAL2416 holding region correctly initiated on second outbound leg.

holding region. However, as shown in Fig. 14, the holding region is correctly reinitialized and updated when the aircraft satisfies the requirements for holding later on the second outbound leg. It should also be noted that even though the reference course for the first holding region was in error, the conflict logic would still detect any aircraft that is predicted to fly through the actual holding pattern.

In most of the cases that we have observed aircraft entering a holding pattern fly generally in the inbound direction and then turn outbound at the holding fix. However, it may be necessary for an aircraft to turn outbound and fly a steady course for a distance before crossing the holding fix and entering the holding pattern. This could cause the holding region to activate prematurely. An example of this case is shown in Fig. 15 (UPS2953). As shown in the figure, the holding region is initialized ($H=1$) as the aircraft settles out on a steady outbound course. However, the logic, described earlier in Section IIID, that prevents the leg length from extending beyond 25 nmi, deactivates the holding region at the $H=0$ point as shown in Fig. 15. Once the holding region has been deactivated ($H=0$ has been set) the 3 requirements for holding (Section IIIB) must be satisfied again in order for the holding region to reinitialize. Figure 16 shows the holding region correctly computed and updated as the aircraft (UPS2953) enters the holding pattern at the fix CADES. Note that in this case the aircraft crosses the holding fix, makes about a 210 deg left turn, then crosses the fix again and enters the pattern with a right turn. This is the so-called tear-drop entry shown in Fig 2, approach (c). The premature activation of a holding region such as that shown in Fig 15 could result in false conflict alerts. A solution to this might be to allow the controller to reinitialize the holding logic for any aircraft at any time through a simple keyboard or trackball input. A re-initialization of the holding logic for any aircraft would then require that the aircraft again satisfy the 3 requirements for holding.




Figure 15. UPS2953 holding region activated prematurely due to long outbound approach to the fix




Figure 16. Holding region for UPS2953 correctly recomputed after entering holding pattern at CADES.

B. Holding Conflict Detection

Figure 17 shows the implementation of the holding function on the CTAS plan-view graphical user interface. One aircraft (CAA558) is holding at FL310 at a fix along the arrival route to DFW, and the other aircraft (FDX3004) is flying level at FL310 along the same arrival route. The 12 sec radar track histories for both aircraft are shown as are the flight data blocks and the dead reckoning trajectory predictions for both aircraft. (The flight data blocks include the aircraft call sign and destination airport in line 1, the flight plan and Mode C altitude in line 2, and the aircraft type and ground speed in line 3.) CTAS computes two trajectories for every aircraft at every track update [5]. The nominal trajectory (not shown in Fig. 17) merges track data and flight plan intent to build a smooth trajectory from current track position back to the intended flight plan. The dead reckoning trajectory (shown in Fig. 17), which is based purely on the track data, extends forward for 2-5 min (5 min in this example). Both trajectories are probed for conflict. The list labeled “Holding Region Conflicts” indicates that a conflict has been detected between the FDX3004 trajectory and the CAA558 holding region. As indicated in the list, FDX3004 will penetrate the holding region in 3 min. In this case, and several other examples that follow, the user has clicked the “show” button on the list to display the trajectories and the holding region.

As shown in Fig. 17, CAA558 is just starting the second turn outbound. Neither Conflict Alert nor conventional Conflict Probe applications would have detected this conflict until possibly the point where FDX3004 had penetrated the holding region and both aircraft were flying nearly head-on. As shown in Fig. 18, FDX3004 was

vectored hard right and then hard left and climbed to FL330 to miss the holding pattern airspace and maintain separation with CAA558.

Note, in Fig. 17, the 5 min dead reckoning trajectory for FDX3004 and the 2 min dead reckoning trajectory for CAA558. The dead reckoning trajectory for any aircraft in holding is reduced from its nominal value of 3-5 min to a 2 min trajectory. Reducing the dead reckoning trajectory to 2 min helps to minimize false alerts caused by dead reckoning trajectories for aircraft in turning flight. But, keeping the dead reckoning trajectory at a low, but non-zero value, serves as a safety net in case the aircraft is intending to fly out of the holding pattern. Once the aircraft flies out of the holding pattern its dead reckoning trajectory is again increased back to its nominal value.

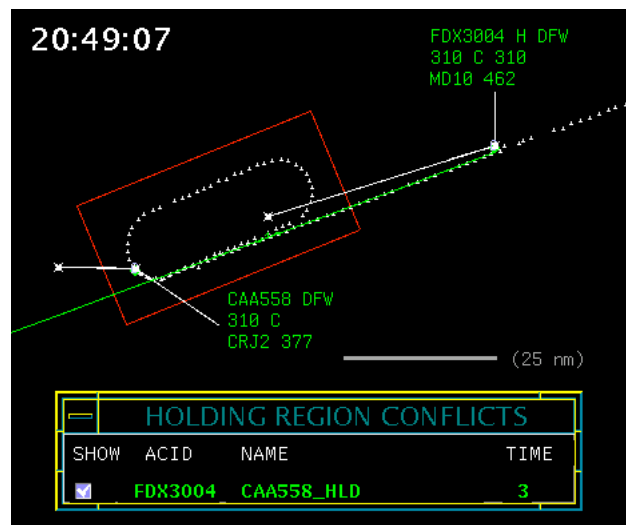


Figure 17. Conflict detected between CAA558 holding at flight plan fix and FDX3004 on flight plan; both level at FL310.

Figure 19 shows a conflict between EGF506, which has just entered the holding pattern at 15,000 ft, and two flights, N718HC and SWA1789, who's trajectories are both descending through the holding pattern. Note from the flight data blocks that N718HC is descending from 15,800 ft to 10,000 ft and SWA1789 is descending from 17,600 ft to 14,000 ft. The outbound (East-North-East) trajectory of EGF506, and the South-bound trajectories of the 2 conflict aircraft are such that the conflict would not have been detected by Conflict Alert or conventional Conflict Probe methods until EGF506 turned inbound towards the holding fix. As shown in the Fig. 20, the crossing traffic was vectored to the left to resolve the conflict with the holding aircraft as it turned inbound towards the holding fix.

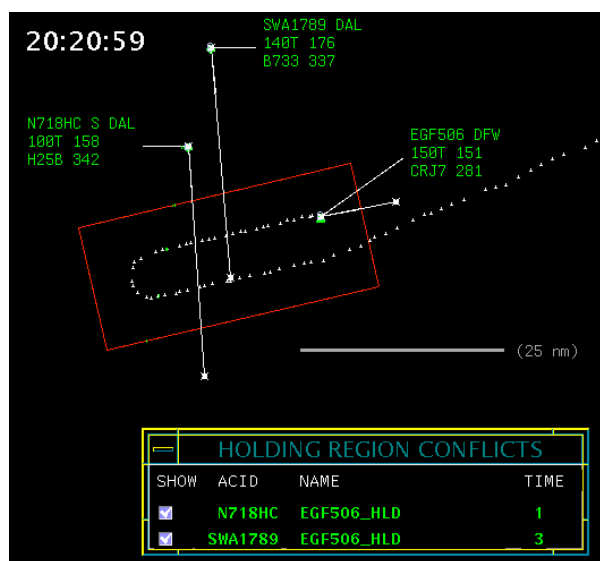


Figure 19. N718HC and SWA1789 crossing holding pattern just prior to EGF506 turning inbound in holding pattern.




Figure 18. FDX3004 vectoring to maintain separation with CAA558 in holding.

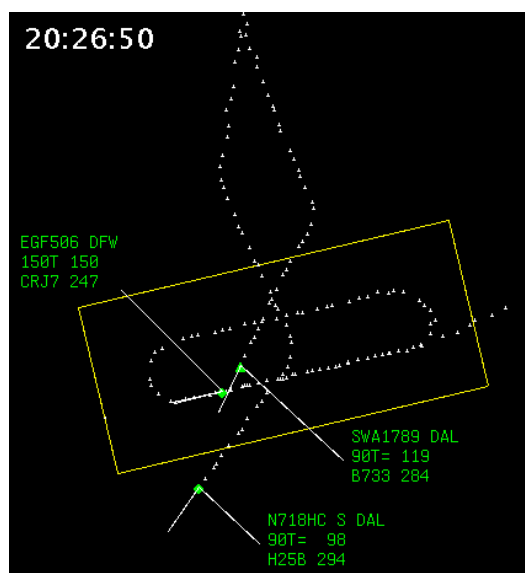


Figure 20. N718HC and SWA1789 vectoring to maintain separation with EGF506.

Shown in Fig. 21 is an example that resulted in an operational error, i.e., the two aircraft passed closer than legally required separation. (This is why the call signs have been changed in this example.) The holding region shown is for AC1, which at this point, had completed several laps around the holding pattern. AC6 was an over-flight aircraft which had turned to the right departing its nominal route of flight (solid line at right near AC6 data block in Fig. 21). The 5 minute dead reckoning trajectory is also shown for AC6. The conflict is detected 3 min prior to the point where AC6 would penetrate the AC1 holding region. The other traffic are included in this example to illustrate the complexity of the traffic situation. Three other aircraft (AC2, AC3, AC4) were holding at the same fix (but at different altitudes) and a level over-flight (AC5) was passing through the pattern at FL330. Note that AC5 flies through the holding region without an alert because no aircraft are holding at FL330. It is important to note that without the dual trajectory probing this conflict would not have been detected until AC6 had been much closer to the holding region. This example resulted in an operational error between AC1 and AC6. The 3 minute warning generated by the holding function would have been plenty of time for the controller to have vectored one of the aircraft and prevented the operational error.

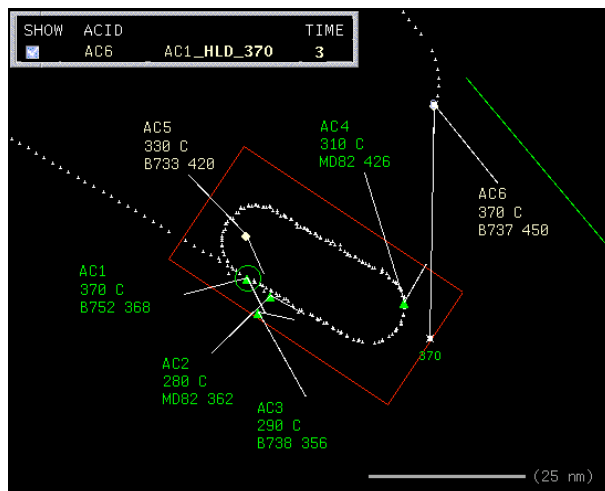


Figure 21. Conflict detected between holding aircraft (AC1) and off-route over-flight aircraft (AC6).

Figure 22 shows a conflict between the dead reckoning trajectory for AAL1212, which has just turned South off of its arrival route, and AAL2416, which has just entered a holding pattern. Both aircraft were level at FL370. In this case AAL1212 was also entering a holding pattern at a fix about 50 nmi to the North-North-East of the pattern in which AAL2416 was flying. This could be considered a false alert since AAL1212 was entering holding and not flying South. Or, it could be a warning to the controller that the holding pattern to the South already had traffic at FL370.

The conflict in Fig. 22 suggests that the time horizon for holding conflict detection might be reduced to something less than 5 min. If the time horizon had been 3 min in this example the conflict would not have been displayed because AAL1212 would have completed the turn outbound (almost due West) before a 3 min trajectory conflicted with the AAL2416 holding region. If, however, AAL1212 had proceeded South-bound, towards the AAL2416 holding pattern, the conflict would have been detected with 3 min to go. As in the Fig. 21 example, 3 min would have been plenty of time for the controller to have resolved any problem that may have developed.

Figure 23 is an example of a false alert. AAL1125 is clearly departing the holding pattern and flying inbound, while the trajectory of the non-holding aircraft, AAL865, is penetrating the holding region at the outbound boundary. While this holding alert would go off as soon as AAL1125 departs the holding region, any false alert is a distraction for a controller. Note the similarity between this example and that in Fig. 17. The relative trajectories of the holding and non-holding aircraft are nearly the same in both examples. The key difference however, is that CAA558 in Fig. 17 is turning outbound for another lap around the pattern and is, therefore, potentially in conflict with FDX3004 flying inbound on the arrival route.

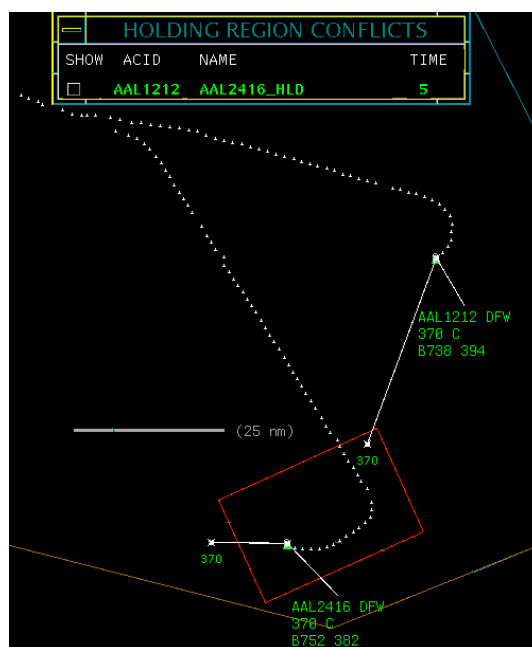


Fig. 22. Conflict example showing dead reckoning trajectories.

Additional logic could distinguish between the false alert in Fig. 23 and the valid alert in Fig. 17. In cases such as this, where both aircraft are flying nearly the same course, the turn status of the holding aircraft could be tested to determine whether or not to display the holding alert. A steady turn status (2-3 consecutive updates) that is consistent with the holding region turn direction would display the alert. A steady no-turn status, and/or, small values of $\Delta\psi_{MF}$ and $\Delta\psi_{NF}$ would not display the alert. A quick controller input that re-initialized the holding logic for any aircraft that has been cleared inbound could also eliminate this type of false alert.

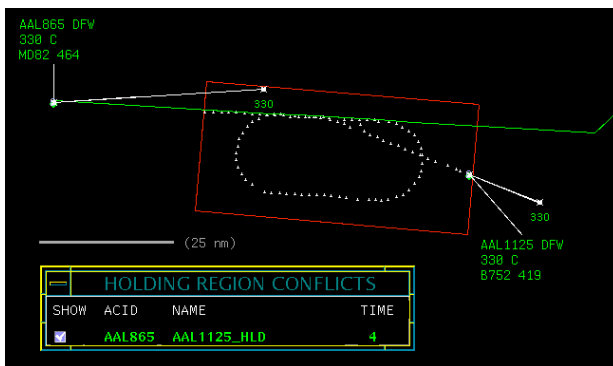


Figure 23. False alert example. Holding aircraft is flying in, while non-holding aircraft conflicts holding region.

VI. User Interface Ideas

A. Display Elements

Figures 17-23 illustrate aspects of a user interface that is envisioned if the holding function were implemented on an air traffic controller's display. Generally there is no holding information displayed unless a trajectory is predicted to conflict with a holding region. The controller might have some way of toggling an active holding region on and off but the holding regions are generally not displayed to the controller once they have been activated. It is anticipated that such a display would be too distracting. When a conflict is predicted the two data blocks could flash as they do under the existing Conflict Alert function. Additionally, there could be an option to toggle the display of the holding region and the conflict aircraft trajectory as shown in Figs. 17-23. In the CTAS PGUI implementation a holding conflict list (Holding Region Conflicts) is displayed and the user toggles the list for a graphic display. The track histories shown in Figs. 17-23 are for illustrative purposes only and would not generally be displayed on a user interface.

B. Manual Activation

There could be circumstances where a controller would want to activate or re-initialize the holding function manually. Manual activation could prevent premature holding initialization for aircraft on long outbound (e.g., Fig. 15). Manual activation could also be used as a means to re-initialize the automatic holding logic for an aircraft that has flown a trajectory that the controller knows could trigger a false hold initialization. A controller could manually activate or re-initialize the holding function by entering the aircraft identifier (e.g., computer identification number) along with some kind of simple "hold-init" entry. Referring to Fig. 15, if a hold-init command were input any time after the point marked "S", a large turn followed by a steady outbound course would be required for initialization and the pattern would correctly initialize as shown in Fig. 16.

The "hold-init" command could give the controller the option of entering the name of the holding fix. The controller might also select a holding fix by using a trial planning tool such as that described in [4]. With the holding fix identified, the point at which the holding region should be initialized, i.e., the H=1 point (see Figs. 4 and 5), must be determined. All of the automatic hold detection logic described above in Section III could be used to compute the H=1 estimate of the holding fix position. When the estimated holding fix position is within a parameter distance (e.g., 3-5 nmi) from the position of the controller selected fix, the holding region is initialized and updated using the known fix and automatic estimates of the turn direction, reference course, and leg length. Recall that if the holding fix is known, the turn radius R and the offset D_O may be computed directly at the H=1 point.

If the controller entered the turn direction along with the holding fix, the large turn would not be required to identify the turn direction, thus further simplifying the initialization logic. It should be emphasized that the results in the previous sections (Sections VA and VB) establish that if the aircraft enters the holding pattern under normal procedures (e.g., fly inbound toward the fix and turn outbound at the fix) these manual activation procedures should not be necessary.

VII. Multiple Aircraft Holding at the Same Fix

It is common for multiple aircraft to hold at the same fix but at different altitudes. It is known from holding operations that aircraft holding at the same fix turn in the same direction and fly a common reference course. This knowledge could be used to improve, or check, our estimate of outbound reference course and holding fix position when a new aircraft enters the holding pattern. In order to develop the logic to handle multiple aircraft in the same pattern it is instructive to first examine multiple holding regions that have been computed independently. Figure 24 shows the radar track histories and the holding regions for 5 aircraft all of which are holding at the fix SJT. SJT is not a flight plan fix for any of the aircraft and so the estimated holding fixes and resulting holding regions are all computed independent of one another. The reference course estimates for the 5 aircraft in the pattern are nearly equivalent (to within about 5 or 10 degrees) and the holding regions are reasonably well correlated with one another. We have observed similar results for other groups of aircraft holding at a common fix.

First note that the estimated position of the holding fix for each aircraft is inside the holding regions for all other aircraft in the pattern. So we can deduce that if the estimated position of the holding fix for a holding candidate (at the H=1 point) is inside the holding region of any aircraft for which holding region has been previously activated, then the two aircraft could be flying the same holding pattern. We would expect the turn direction of the new aircraft to match the turn direction of the existing aircraft in the pattern. We would also expect the reference course of a new aircraft to be close to the reference course of any previously activated holding aircraft.

The correlation of data from multiple holding aircraft could be used to detect a reference course estimate that is substantially in error such as the case shown in Fig. 13. Note that the aircraft in Fig. 13 (AAL2416) is one of the aircraft shown holding in Fig. 24. If the initial estimate of the reference course for AAL2416 were compared against the reference course of aircraft already in the pattern, it could have been identified as being uncorrelated by about 30 deg. In such a case the initialization of the holding region could be delayed until more outbound track data were observed and processed by the tracking filter. Alternatively, if there were enough confidence in the reference headings of those aircraft already in the pattern, the initial estimate for a new aircraft could be replaced with the average of the reference headings for aircraft already in the pattern. The data shown in Fig. 24 suggest a high confidence level in the actual holding pattern course being very close to the average of the reference headings that were automatically computed for the 5 aircraft (or 4 aircraft) shown. From the position of the holding regions relative to the track data it is noted that the automatically computed holding fix positions for all 5 aircraft are reasonably close (e.g., within a few miles) to one another.

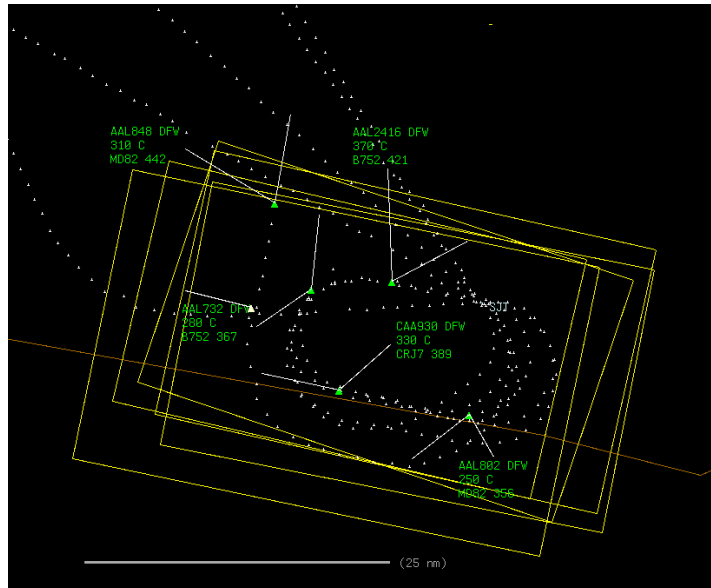


Figure 24. Holding regions for 5 aircraft holding at the same off-route fix (SJT).

VIII. Concluding Remarks

A holding function that automatically detects aircraft that have entered a holding pattern, computes and updates a model of holding pattern airspace, and probes the trajectories of other aircraft for conflict with the holding airspace at the holding altitude has been developed, implemented in software, and tested with holding aircraft data from the Fort Worth Air Route Traffic Control Center.

The results show that real-time processing of Center Host radar track and flight plan data alone may be used to automatically identify holding aircraft and model holding pattern airspace.

Of 37 aircraft that went into holding during a 1-hour period in Fort Worth Center, 34 holding regions were automatically computed and updated correctly. Three were initially computed incorrectly and then later corrected.

A holding region could activate incorrectly or prematurely if the aircraft flies an extended outbound leg before entering the holding pattern, or if it enters the pattern via a course that includes off-nominal turns. This occurred for 3 out of the 37 holding aircraft. In all cases the holding region was automatically deactivated after several track updates and then reactivated correctly on the second outbound holding leg.

Several example cases from recorded Fort Worth Center data were shown where the holding function detected a potential traffic conflict that would not have been identified by either Conflict Alert or Conflict Probe automation. In one example, which resulted in an operational error, the holding function detected the conflict 3 minutes prior to loss of separation. This would have been sufficient time for the controller to have vectored one of the aircraft and prevented the operational error.

A simple user input was proposed that would allow the controller to re-initialize the holding logic at any time. This would reduce the chance of false holding conflict alerts.

IX. References

1. FAA Air Traffic Publications Library, Aeronautical Information Manual, Chapter 5 Air Traffic Procedures, Section 5-3-7 Holding. Effective Feb 21, 2002. See also <http://www2.faa.gov/atpubs/AIM/>.
2. Dr. Tsuyoshi Goka, John E. Robinson III, and Michelle M. Eshow, "Course Estimation and Turn Detection Algorithms for the Center-TRACON Automation System (CTAS)" Draft August, 2004.
3. Denery, D. G. and Erzberger, H.: "The Center/TRACON Automation System: Simulation and Field Testing," Proceedings of the advanced Workshop on ATM (ATM 95) sponsored by the National Research Council of Italy; Oct. 2-6, 1995; Capri Italy; Also published as a NASA Technical Memorandum 110366, August, 1995. See also <http://www.ctas.arc.nasa.gov/>
4. McNally, et al, "Operational Evaluation of the Direct-To Controller Tool" 4th USA/Europe ATM R&D Seminar, Santa Fe, New Mexico. Dec 3-7, 2001. See also <http://www.ctas.arc.nasa.gov/>
5. Erzberger, H, "Transforming the NAS: The Next Generation Air Traffic Control System" NASA/TP-2004-212828, June 2004.